

Mobility Map Computations for Autonomous Navigation using an RGBD Sensor

Nicolò Genesio[†], Tariq Abuhashim[†], Fabio Solari[‡], Manuela Chessa[‡] and Lorenzo Natale[†]

Abstract—In recent years, the numbers of life-size humanoids as well as their mobile capabilities have steadily grown. Stable walking motion and control for humanoid robots are active fields of research. In this scenario an open question is how to model and analyse the scene so that a motion planning algorithm can generate an appropriate walking pattern.

This paper presents the current work towards scene modelling and understanding, using an RGBD sensor. The main objective is to provide the humanoid robot iCub with capabilities to navigate safely and interact with various parts of the environment. In this sense we address the problem of traversability analysis of the scene, focusing on classification of point clouds as a function of mobility, and hence walking safety.

I. INTRODUCTION

The potential market of service and entertainment humanoid robots has attracted great research interests. One of the most fundamental and challenging steps is to allow robots to interact and walk autonomously within a real world scenario. Ideally, they should be able to accept high level human commands and, then, autonomously walk in a real-life environment consisting of floors and stairs without colliding with obstacles. Thus, it is necessary for humanoids to be able to identify candidate traversable places within acceptable stability limits. This includes, for instance, flat surfaces or surfaces with limited slope and roughness values.

This paper presents a pipeline to compute a mobility map in real-time using an RGBD sensor mounted on the iCub, as shown in Figure 1. As in most navigation scenarios the terrain information is vital to classify what is traversable and what is not. In particular, legged-type robots such as humanoids require precise information on the surrounding terrain, not only for determining locomotion strategies, but also for coordinating complex body motions, such as generating obstacle-avoiding free-leg trajectories during walking.

This research has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement No. 611909 (KoroiBot).

[†] are with the iCub Facility, Istituto Italiano di Tecnologia, via Morego, 30, 16163 Genova, Italy email: {nicolo.genesio, tariq.abuhashim, lorenzo.natale}@iit.it

[‡] are with the DIBRIS, Università degli studi di Genova, Via All'Opera Pia, 13, 16145 Genova, Italy email: {fabio.solari, manuela.chessa}@unige.it




Fig. 1: Example of what we want to obtain, a map where each region has a score between 0(red) and 1(green) that indicates how much traversable a surface is.

For humanoids, and from a simplified kinematic point of view, a surface is considered as a part of an obstacle if its geometric properties do not allow for a safe foot step planning. Useful surface properties may include its slope, roughness, rugosity and size. Such properties can be measured or estimated from 3D data, which can be provided, for instance, by using a structured light sensor. Such sensors are considered cheaper and lighter than their counterpart laser scanners. They also provide richer visual information, which can be useful for tasks including object recognition, classification and segmentation.

II. RELATED WORK

In recent years, the interest in scene modelling and understanding algorithms has widely spread. Specifically, segmentation and classification of outdoor workspaces using 3D depth data [1], in order to allow mobile robots to move in real world environments while navigating between obstacles safely. [2].

The literature is rich also of research examples that deals with the problem of indoor scene understanding in order to give to the humanoid robot the capability of moving in autonomy in a domestic environment [3]. The terrain plays a key role in works like this, and an overview on terrain traversability methods can be found in [4]. Most of the methods use 3D geometric features (*e.g.* normals) to interpret the scene that in general are very computationally demanding. For this reason these approaches are not appropriate to manage dynamic obstacles, for which the map has to be updated frequently and then the computation has to be repeated several times. On the other hand approaches that *learn* the scene through proprioceptive

informations are more suitable for online applications. This approach is demonstrated in [5] for a Mars-Rover platform where the authors represent the environment in proprioception space in terms of expected slip. Moreover proprioceptive informations can be used in addition of 3D data to increase the reliability of the slope inference and then the accuracy of the classification [6].

More in general the reliability in obstacle detection and scene classification can be augmented relying on multiple sensing modalities such as color and 3D data [7]. Visual and structural modalities are clearly complementary: vision alone may be inadequate or unreliable in the presence of strong shadows, while depth measurement of sensors like structured light sensors can be misled by the sunlight interference with the IR sensor. In literature can be found works that try to solve the same problem, but using different tools to obtain this kind of information as laser range-finder [7], or stereo vision [8], but vision techniques are in general challenged by darkness or other extreme lighting conditions and the accuracy that they can provide is too low to manage a such complex issue as scene modelling and interpretation.

III. OUR IMPLEMENTATION

This section describes our mobility map implementation, where Figure 2 highlights different data processing and estimation steps.

Data reduction: In this paper, we used the Asus Xtion PRO LIVE. This sensor produces depth maps at 30 fps with a resolution of 640×480 . The depth map includes over 300 thousands of points. In the next steps, it is required to evaluate a surface properties at any point of interest using its k -nearest neighbours. The computation cost is, hence, $\mathcal{O}(nk)$, where n is the number of points in the depth map. Based on the application, the way the sensor is mounted on the robot and the sensor field of view, we are able to reduce the computational cost of the depth map by downsampling it into voxels of $1\text{cm} \times 1\text{cm} \times 1\text{cm}$. Further data reduction can be achieved by limiting the depth range of the sensor to include only points with acceptable depth uncertainty. In this paper, we removed points at distances longer than 1.5-2 meters.

Denoising: 3D information contained in the point cloud is often contaminated by noise. This is due to many causes including the interference of daylight with the IR sensor and intrinsic errors made by the sensor performing the triangulation. Before attempting to estimate the characteristics of a point with respect to its surrounding, it is important to analyse if the surrounding neighbourhood is a good representation of the underlying sampled surface. Thus, following the method proposed in [9], for each point p_q in a cloud P , the mean distance \bar{d} to its k -nearest neighbours is first

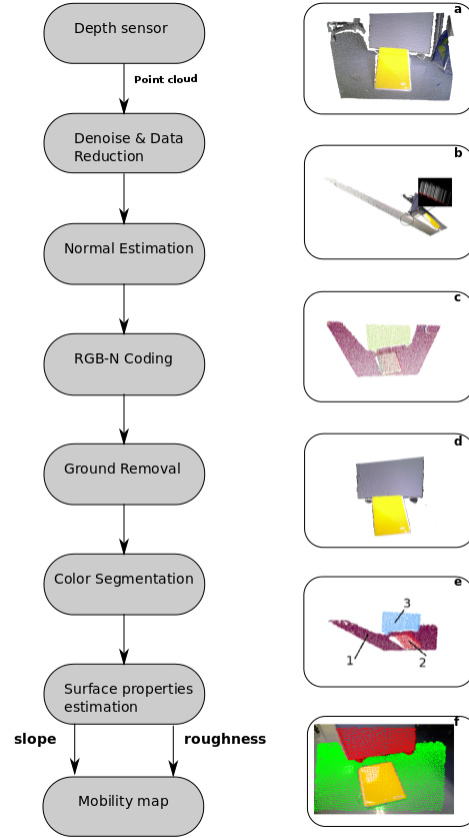


Fig. 2: Block diagrams displaying the main points of our algorithm.

computed. Then a distribution over the mean distance space for the entire point cloud is assembled and its mean μ_k and standard deviation σ_k are estimated. Our motivation is to keep points with mean distance \bar{d} to the k -nearest neighbours that is statistically similar to those by rest of the points. Thus, the new down-sampled point cloud P^* can be extracted using,

$$P^* = \{p_q^* \in P | (\mu_k - \alpha\sigma_k) \leq \bar{d} \leq (\mu_k + \alpha\sigma_k)\} \quad (1)$$

where α is a desired density restrictiveness factor.

Normal Computation: For a good representation and segmentation we have to use features with high discriminating power such as normals and curvature. They are two of the most widely used geometric features because they provide information on the orientation of surfaces, indispensable for scene understanding. They are treated as local features, because they characterize the information provided by the k -nearest neighbours of each point. Their estimated values are sensitive to sensor noise and the selection of the k neighbours.

After we have determined the neighbourhood P^k of a query point p_q , we can use it to compute a local feature that represents the geometry around the query point. One surface point feature can be computed as

the normal vector of the tangent plane, which can be estimated solving a least-square plane fitting problem over P^k [10]. As explained in [11] this plane can be represented by a point x and a normal vector \vec{n} , and, given the distance from a point $p_i \in P^k$ to the plane as $d_i = (p_i - x) \cdot \vec{n}$, the values of x and \vec{n} are computed in a least-square sense such that $d_i = 0$.

Assume

$$x = \bar{p} = \frac{1}{k} \sum_{i=1}^k p_i \quad (2)$$

as centroid of P^k , we can then solve the fitting problem obtaining \vec{n} as the eigenvector \vec{v}_0 of the smallest eigenvalue λ_0 of the covariance matrix $C \in R^{3 \times 3}$ of P^k , expressed as

$$C = \frac{1}{k} \sum_{i=1}^k (p_i - \bar{p}) \cdot (p_i - \bar{p})^T. \quad (3)$$

To resolve for the sign of \vec{n} , however, we need the viewing point of the sensor v_p . To define the sign of the normals, each \vec{n}_i has to satisfy

$$\vec{n}_i \cdot (v_p - p_i) > 0. \quad (4)$$

In this way all the normals \vec{n}_i point towards the viewpoint of the sensor.

RGB-N coding: Before segmentation, we colour-code the point cloud using the computed normals. In the previous step, we have computed, for each point p_i , a normal $\vec{n}_i = \{n_{ix}, n_{iy}, n_{iz}\}$, function of its neighbourhood. We assign to p_i the following RGB-coding

$$\begin{aligned} R &= \frac{255}{\max_{nx} - \min_{nx}} * (n_{ix} - \min_{nx}) \\ G &= \frac{255}{\max_{ny} - \min_{ny}} * (n_{iy} - \min_{ny}) \\ B &= \frac{255}{\max_{nz} - \min_{nz}} * (n_{iz} - \min_{nz}), \end{aligned} \quad (5)$$

where \max_{n*} and \min_{n*} are respectively the maximum and minimum value of the component $n*$ among all the normals contained in the cloud. By doing so, point-cloud's surfaces with similar normal, and thus orientation, will have consistent and similar colors (Fig. 2.c). This is an important step which allows us to use color differences to segment regions with common geometric features.

Ground removal: Ground removal is an important step for robust segmentation results. It provides two main benefits: First, it has the effect of isolating the object from the background, and thus improving segmentation. Second, it allows us to further reduce the computational cost of the entire system, because we exclude from the analysis the segment that has the majority of the points in the cloud.

There are many ways to segment the floor from the rest of the scene, some of them use images [12] but,

despite their robustness, they are often afflicted by the differences in light conditions of the scene. One way to achieve robustness is by using RANSAC, using the following steps:

- 1) Randomly select three non-collinear unique points $\{p_i, p_j, p_k\}$ from the point cloud P ;
- 2) Compute the model coefficients from the three points ($ax + by + cz + d = 0$);
- 3) Compute the distances from all $p \in P$ to the plane model (a, b, c, d) ;
- 4) Count the number of points $p^* \in P$ whose distance d to the plane model falls between $0 \leq |d| \leq |d_t|$, where d_t represents a user specified threshold.

The last step represents a way of “scoring” a specific model that we used to find the best plane in the cloud. Every set of points p^* is stored, and the above steps are repeated for k iterations. The number of iterations is defined as follow. If ϵ is the probability of picking a sample that produces a bad estimate (*i.e.* outlier), then $1 - \epsilon$ is the probability of picking at least one good sample (*i.e.* inlier). This means that the probability of picking γ good samples becomes $(1 - \epsilon)^\gamma$. For k trials, the probability of failure becomes $(1 - (1 - \epsilon)^\gamma)^k$. If p is the desired probability of success (*e.g.* $p = 0.99$), then:

$$1 - p = (1 - (1 - \epsilon)^\gamma)^k \implies k = \frac{\log(1 - p)}{\log(1 - (1 - \epsilon)^\gamma)}. \quad (6)$$

After the algorithm is terminated, the set with the largest number of points (inliers) is selected as the support for the best planar model found, the ground in our case.

Color segmentation: After removing the ground, our objective is to segment the remaining regions in the cloud. A RANSAC-based method won't suffice for this task, because it assumes that all the objects in the scene can be mathematically modelled. Thus, we utilize *Region Growing* [13] for this task. This method follows a *flood fill* approach. The method aims at selecting a set of *homogeneous* points by optimising for a given inter-regional constraint. A *seed* point is first selected. At each optimisation step, the surrounding of the seed point is iteratively allowed to *grow* by including more points into the computations of the local constraint. This process is iterated until this local constraint is satisfied. Typical constraints may include the Euclidean distance between the selected point and the seed point or local features, including geometric and photometric features. For more efficient computations, we utilise the previously color-coded point-cloud based on the computed surface normals as a feature, and formulate our regional constraints using the Euclidean distance. This choice simply allows us to extract regions with homogeneous surface

```

while  $\{A\}$  is not empty do
  Current Region  $\{R_c\} \leftarrow \emptyset$ ;
  Current seeds  $\{S_c\} \leftarrow \emptyset$ ;
  Select randomly  $p$  from  $\{A\}$ ;
   $\{S_c\} \leftarrow \{S_c\} \cup p$ ;  $\{R_c\} \leftarrow \{R_c\} \cup p$ ;
   $\{A\} \leftarrow \{A\} \setminus p$ ;
  for  $i=0$  to  $\text{size}(\{S_c\})$  do
    Nearest neighbours of  $p$ :  $B_c \leftarrow \Omega(\{S_c\})$ ;
    for  $j=0$  to  $\text{size}(\{B_c\})$  do
      Current neighbour point  $p_j$ ;
      if  $\{A\}$  contains  $p_j$  and
       $\text{coldist}_{pj} < c_{th}$  then
         $\{S_c\} \leftarrow \{S_c\} \cup p_j$ ;
         $\{R_c\} \leftarrow \{R_c\} \cup p_j$ ;
         $\{A\} \leftarrow \{A\} \setminus p_j$ ;
      end
    end
  end
   $\{R\} \leftarrow \{R\} \cup \{R_c\}$ ;
end

```

Algorithm 1: Region Growing based on color coding of surface normals.

orientations. Algorithm 1 summarizes steps of Region Growing on 3D point-clouds, where P is a color-coded point-cloud, c_{th} is a color threshold, coldist_{pj} is the computed Euclidean distance between points p and p_j in the color space ($Y C_r C_b$), R is the region list, $\Omega(\cdot)$ is the neighbour finding function, and A is the available point list, that it is initialized using all the points of P .

Surface properties estimation: After dividing the cloud in clusters with similar geometric properties (Fig. 2.e), we estimate those properties that will be used to determine the mobility scoring of individual segments in the point-cloud. First, we compute the slope, which is a measure of steepness for planes or more in general for flat regions. First, we estimate the best plane that fits the points within the i -th segment using least-squares and RANSAC. The slope s_i is then calculated as the angle between the fitted plane normal \vec{n}_i and the vector \vec{n}_f representing the normal vector of the removed ground segment. This can be simply computed as:

$$\theta = \arccos\left(\frac{\vec{n}_i \cdot \vec{n}_f}{\|\vec{n}_i\| \cdot \|\vec{n}_f\|}\right) * \frac{180}{\pi}. \quad (7)$$

In addition to the slope, our mobility function considers the surface roughness, which represents as a measure of the asperity of a certain surface. Surface roughness is widely used in mobile robotics navigation planning [14], because of its ability to limit the mobility of various robotic platforms. For humanoids, safe and stable foot step planning is highly correlated to terrain roughness.

There are many ways to estimate the roughness mostly because a canonical definition has not already formulated. We followed an approach already used in [15]. For each segment of the cloud we compute the roughness index r as:

$$r = \frac{A}{A'}, \quad (8)$$

where A is the area of the segment, and A' is the area of the segment projection on the corresponding estimated plane, which was computed using RANSAC and least-squares. To compute r we utilize *Delaunay triangulation*. Given a set of vertices $V = \{v_m\}$, such that $v_m \in \mathbb{R}^3$ and $m = 1, \dots, MS$, and $v_m = (x_m, y_m, z_m)$ represents the vertex m described by its x, y, z point coordinates. The triangles of the surface are contained in the set $T = \{t_n\}$, where $n = 1, \dots, N$, such that $t_n \subset V$ and $t_n = (v_{1_n}, v_{2_n}, v_{3_n})$ represents a triangle defined by three vertices in V . Thus, A , and similarly A' , are computed using:

$$A = \sum_{j=1}^N a_j, \quad (9)$$

where a_j is the area of the j -th triangle (t_j), which is computed as half the magnitude of the cross product of the vectors $(v_1 \vec{v}_{2_j}$ and $v_2 \vec{v}_{3_j})$ representing two adjacent sides of the triangle. Thus,

$$a_j = \frac{1}{2} \|\vec{v}_1 \vec{v}_{2_j} \times \vec{v}_2 \vec{v}_{3_j}\|. \quad (10)$$

Mobility map: The classification of the scene as a function of robots mobility is performed using a simple and intuitive approach. Every segment is assigned a score y between 0 and 1, where 0 corresponds to *untraversable* surfaces while 1 corresponds to *traversable* surfaces that does not require any additional foot step planning, which is shown in table I. This mobility rule is very discrete, and hence limits its applications to multiple foot step configurations. In order to accommodate finer mobility decisions, we further improve the resolution of our decisions on the traversability of surfaces by applying Gaussian Process Regression (GPR) with a Squared Exponential (SE) covariance function [16] (see Fig. 3):

$$k_y(x_i, x_j) = \sigma_f^2 \exp\left(-\frac{1}{2\lambda^2}(x_i - x_j)^2\right), \quad (11)$$

where (x_i, x_j) are input pairs containing surface features, namely the estimated slope and roughness, $x_i = (s_i, r_i)^\top$, σ_f is the variance (we assumed, $\sigma_f = 1$) and λ is the length scale which defined the smoothness of the mobility score (we assumed, $\lambda = 0.1$). Since we would like our mobility to be a function of both slope and roughness, we train our GPR model using two the inputs (s and r) with one output mobility score m .

Thus, we can obtain the score as an inference instance using our trained GPR model:

$$\hat{m}(x_t) = \mathbf{k}_*^\top (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y}, \quad (12)$$

where, \mathbf{K} is the covariance matrix, \mathbf{k}_* to denote the vector of covariances between the test point $x_t = (s_t, r_t)^\top$ and the training points $x_i = (s_i, r_i)^\top$, σ_n is the expected noise in the measured mobility \mathbf{y} , given the normals. The range of input dimensions was limited by the walking capabilities and mechanical limits of the robot.

$s \uparrow \rightarrow$ $r \downarrow$	0-10	10-20	20-30	30-40	>40
1-1.2	1	0.75	0.5	0.25	0
1.2-1.4	0.75	0.5	0.25	0	0
1.4-1.6	0.5	0.25	0	0	0
1.6-1.8	0.25	0	0	0	0
1.8-2.0	0	0	0	0	0
2.0-2.2	0	0	0	0	0
>2.2	0	0	0	0	0

TABLE I: Mobility rule estimated empirically.




Fig. 3: A sample mobility cost function, which was generated with GPR.

IV. EXPERIMENTS

In this Section we present supporting experimental results to demonstrate the performance of the proposed mobility map computation pipeline. The collected dataset comprised 3D depth maps and associated RGB images. Due to the limited mobility of the iCub robot and the sensitivity of the depth sensor to daylight, we tested our implementation only on data sets collected for indoor environments. For example, laboratory corridors with various objects being located along the way. The method of this paper was implemented using C++, and is available on GitHub¹.

¹<https://github.com/Nicogene/MobilityMapBuilder>.

A. RGBD Sensor Model

The validation of our system has been made through the projection of the 3D mobility maps on the images acquired by the RGB camera of Asus Xtion Pro Live (Fig. 4, 6). The mapping has been done using the following pinhole camera model,

$$\begin{aligned} x &= c_x + f_x * \frac{X}{Z} - o_x \\ y &= c_y + f_y * \frac{Y}{Z} - o_y, \end{aligned} \quad (13)$$

where $\{x, y\}$ are the image coordinate, $\{X, Y, Z\}$ are the 3D point coordinates, $\{c_x, c_y\}$ define the optical center, $\{f_x, f_y\}$ are the focal lengths and $\{o_x, o_y\}$ are the factory offsets between the IR sensor frame and the RGB camera frame.

B. Free corridor




Fig. 4: Images displaying the path crossed, in sequence from (a) to (f). We projected on them the clusters coloured in function of the traversability; we used a color map where green means *mobility score*=1 and red means *mobility score*=0.

Fig. 4 shows mobility map estimation and point-cloud segmentation results in the image frame using our pipeline for the case of obstacle free path. Our motivation is to initially test our implementation of ground removal, since the performance of our method relies on it. Also, showing the results in the image frame provides us with a way to validate the accuracy of our implementation. On the other hand, Fig. 5




Fig. 5: Figures presenting point clouds acquired and segmented with our normals/color based method.

shows the verification of our segmentation results in 3D space.

C. Corridor with obstacles

Here, we present the results of the usage of our system in conditions more challenging than the one we talked about previously. Fig. 6 shows the projections of the computed mobility maps into the images, while Fig 7 shows the segmented point-clouds data.

D. Processing Requirements

The dataset was processed on an Intel Core i5 CPU@2.30GHz×4. In table II we report the time required for each step of the algorithm measured on a sample cloud. Notice that the variations in the number of points processed at each step is due to modifications to the point-cloud done by the previous step, including sampling, outliers removal, and ground separation. This analysis demonstrates that the algorithm can be used to analyse a scene at the rate of two frames per second, which is suitable for foot step planning.

V. CONCLUSION AND FUTURE WORK

This paper has presented an approach for scene segmentation and mobility estimation based on the analysis of 3D data acquired by a RGBD sensor. The method was used to successfully and efficiently classify real world indoor workspaces into traversable




Fig. 6: Images displaying the path crossed, in sequence from (a) to (f). We projected on them the clusters coloured in function of the traversability; we used a color map where green means *mobility score*=1 and red means *mobility score*=0.

Step	Time(ms)	Points processed
Denoise	68.133	307200
Data reduction	39.873	307200
Normal estimation	144.766	12004
RGB-N coding	0.203	12004
Ground removal	31.482	12004
Color segmentation	24.961	887
Surface properties estimation	175.303	887
Mobility mapping	1.328	887
Total:	486.049	

TABLE II: Time performance of our algorithm for analysing one sample point cloud.

and untraversable regions. We have achieved a mobility map computation pipeline that runs at 2Hz. Our future work will investigate removing many existing assumptions, including our ground removal algorithm which assumes that the largest segment in any depth map corresponds to the ground. We will also be looking at integrating vision with the depth data to build more robust features for segmentation.

REFERENCES

- [1] C. Dima, N. Vandapel, and M. Hebert, "Classifier Fusion for Outdoor Obstacle Detection," in *International Conference on Robotics and Automation*, vol. 1. IEEE, April 2004, pp. 665 – 671.




Fig. 7: Figures presenting point clouds acquired and segmented with our normals/color based method.

- [2] F. Ferri, M. Gianni, M. Menna, and F. Pirri, "Point Cloud Segmentation and 3D Path Planning for Tracked Vehicles in Cluttered and Dynamic Environments," in *In Proceedings of the 3rd IROS Workshop on Robots in Clutter: Perception and Interaction in Clutter*, 2014.
- [3] T.-Y. Li, P.-F. Chen, and P.-Z. Huang, "Motion planning for humanoid walking in a layered environment," in *Robotics and Automation, 2003. Proceedings. ICRA '03. IEEE International Conference on*, vol. 3, Sept 2003, pp. 3421–3427 vol.3.
- [4] P. Papadakis, "Terrain traversability analysis methods for unmanned ground vehicles: A survey," *Engineering Applications of Artificial Intelligence*, vol. 26, no. 4, pp. 1373 – 1385, 2013.
- [5] A. Angelova, L. Matthies, D. Helmick, and P. Perona, "Learning and prediction of slip from visual information," *Journal of Field Robotics*, vol. 24, no. 3, pp. 205–231, 2007.
- [6] S. Karumanchi, T. Allen, T. Bailey, and S. Scheding, "Non-parametric Learning to Aid Path Planning over Slopes," *I. J. Robotic Res.*, vol. 29, no. 8, pp. 997–1018, 2010.
- [7] C. Rasmussen, "Laser Range-, Color-, and Texture-based Classifiers for Segmenting Marginal Roads," in *IEEE International Conference on Computer Vision and Pattern Recognition Technical Sketches*, 2001.
- [8] Y. Takaoka, Y. Kida, S. Kagami, H. Mizoguchi, and T. Kanade, "3d map building for a humanoid robot by using visual odometry," in *Systems, Man and Cybernetics, 2004 IEEE International Conference on*, vol. 5, Oct 2004, pp. 4444–4449 vol.5.
- [9] R. B. Rusu, Z. C. Marton, N. Blodow, M. Dolha, and M. Beetz, "Towards 3D Point Cloud Based Object Maps for Household Environments," *Robotics and Autonomous Systems Journal (Special Issue on Semantic Knowledge in Robotics)*, vol. 56, no. 11, pp. 927–941, 30 November 2008.
- [10] C. M. Shakerji *et al.*, "Least-squares fitting algorithms of the nist algorithm testing system," *Journal of Research-National Institute of Standards and Technology*, vol. 103, pp. 633–641, 1998.
- [11] R. B. Rusu, "Semantic 3D Object Maps for Everyday Manipulation in Human Living Environments," Ph.D. dissertation, Technische Universität München, 2009.
- [12] S. Kumar, A. Dewan, and K. M. Krishna, "A Bayes Filter Based Adaptive Floor Segmentation with Homography and Appearance Cues," in *Proceedings of the Eighth Indian Conference on Computer Vision, Graphics and Image Processing*,

ser. ICVGIP '12. New York, NY, USA: ACM, 2012, pp. 54:1–54:8.

- [13] R. Adams and L. Bischof, "Seeded region growing," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 16, no. 6, pp. 641–647, June 1994. [Online]. Available: <http://dx.doi.org/10.1109/34.295913>
- [14] D. Stavens and S. Thrun, "A Self-supervised Terrain Roughness Estimator for Off-Road Autonomous Driving," in *In Proc. of Conf. on Uncertainty in AI (UAI)*, 2006, pp. 13–16.
- [15] A. Friedman, O. Pizarro, S. B. Williams, and M. Johnson-Roberson, "Multi-Scale Measures of Rugosity, Slope and Aspect from Bathymetric Stereo Image Reconstructions," *PLoS ONE*, vol. 7, no. 12, p. e50440, 12 2012.
- [16] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.